

Optimal Sampling-Based Motion Planning under Differential Constraints: the Drift Case with Linear Affine Dynamics

Edward Schmerling

Lucas Janson

Marco Pavone

Abstract—In this paper we provide a thorough, rigorous theoretical framework to assess optimality guarantees of sampling-based algorithms for drift control systems: systems that, loosely speaking, can not stop instantaneously due to momentum. We exploit this framework to design and analyze a sampling-based algorithm (the Differential Fast Marching Tree algorithm) that is asymptotically optimal, that is, it is guaranteed to converge, as the number of samples increases, to an optimal solution. In addition, our approach allows us to provide concrete bounds on the rate of this convergence. The focus of this paper is on mixed time/control energy cost functions and on linear affine dynamical systems, which encompass a range of models of interest to applications (e.g., double-integrators) and represent a necessary step to design, via successive linearization, sampling-based and provably-correct algorithms for non-linear drift control systems. Our analysis relies on an original perturbation analysis for two-point boundary value problems, which could be of independent interest.

I. INTRODUCTION

A key problem in robotics is how to compute an obstacle-free and dynamically-feasible trajectory that a robot can execute [1]. The problem, in the simplest setting where the robot does not have kinematic/dynamical (in short, differential) constraints on its motion and the problem becomes one of finding an obstacle-free “geometric” path, is reasonably well-understood and sound algorithms exist for most practical scenarios. However, robotic systems *do* have differential constraints (e.g., momentum), which most often cannot be neglected. Despite the long history of robotic motion planning, the inclusion of differential constraints in the planning process is currently considered an open challenge [2], in particular with respect to guarantees on the quality of the obtained solution and class of dynamical systems that can be addressed. Arguably, the most common approach in this regard is a decoupling approach, whereby the problem is decomposed in steps of computing a collision-free path (neglecting the differential constraints), smoothing the path to satisfy the motion constraints, and finally reparameterizing the trajectory so that the robot can execute it [2]. This approach, while oftentimes fairly computationally efficient, presents a number of disadvantages, including computation of trajectories whose cost (e.g., length or control effort) is far from the theoretical optimum or even failure in finding any solution trajectory due to the decoupling scheme itself [2]. For these reasons, it has been advocated that there is

a need for planning algorithms that solve the differentially-constrained motion planning problem (henceforth referred to as the DMP problem) *in one shot*, i.e., without decoupling.

Broadly speaking, the DMP problem can be divided into two categories: (i) DMP for driftless systems, and (ii) DMP for drift systems. Intuitively, systems with drift constraints are systems where from some states it is impossible to stop instantaneously (this is typically due to momentum). More rigorously, a system $\dot{\mathbf{x}} = f(\mathbf{x}, \mathbf{u})$ is a drift system if for some state \mathbf{x} there does not exist any admissible control \mathbf{u} such that $f(\mathbf{x}, \mathbf{u}) = 0$ [1]. For example the basic, yet representative, double integrator system $\ddot{\mathbf{x}} = \mathbf{u}$ (modeling the motion of a point mass under controlled acceleration) is a drift system. From a planning perspective, DMP for drift systems is more challenging than its driftless counterpart, due, for example, to the inherent lack of symmetries in the dynamics and the presence of regions of inevitable collision (that is, sets of states from which obstacle collision will eventually occur, regardless of applied controls) [1].

To date, the state of the art for one-shot solutions to the DMP problem (both for driftless and drift systems) is represented by sampling-based techniques, whereby an explicit construction of the configuration space is avoided and the configuration space is probabilistically “probed” with a sampling scheme [1]. Arguably, the most successful algorithm for DMP to date is the rapidly-exploring random tree algorithm (RRT) [3], which incrementally builds a tree of trajectories by randomly sampling points in the configuration space. However, the RRT algorithm lacks optimality guarantees, in the sense that one can prove that the cost of the solution returned by RRT converges to a suboptimal cost as the number of sampled points goes to infinity, almost surely [4]. An asymptotically-optimal version of RRT for the geometric (i.e., without differential constraints) case has been recently presented in [4]. This version, named RRT*, essentially adds a rewiring stage to the RRT algorithm to counteract its greediness in exploring the configuration space. Prompted by this result, a number of works have proposed extensions of RRT* to the DMP problem [5, 6, 7, 8, 9], with the goal of retaining the asymptotic optimality property of RRT*. Care must be taken in arguing optimality for drift systems in particular, as the control asymmetry requires a consideration of both forward-reachable and backward-reachable trajectory approximations. Even in the driftless case, the matter of assessing optimality is quite subtle, and hinges upon a careful characterization of a system’s locally reachable sets in order to ensure that a planning algorithm examines “enough volume” in its operation, and thus enough sample points, to ensure asymptotic optimality [10]. Another approach to asymptotically optimal DMP planning is given by STABLE SPARSE RRT which achieves optimality through random control propagation instead of connecting

Edward Schmerling is with the Institute for Computational & Mathematical Engineering, Stanford University, Stanford, CA 94305, schmerling@stanford.edu.

Lucas Janson is with the Department of Statistics, Stanford University, Stanford, CA 94305, ljanson@stanford.edu.

Marco Pavone is with the Department of Aeronautics and Astronautics, Stanford University, Stanford, CA 94305, pavone@stanford.edu.

This work was supported by an Early Career Faculty grant from NASA’s Space Technology Research Grants Program (Grant NNX12AQ43G).

sampled points using a steering subroutine [11]. This paper, like the RRT* variations, is based on a steering function, although it may be considered less general, as it is our view that leveraging as much knowledge as possible of the differential constraints while planning is necessary for the goal of planning in real-time. In our related work [10] we provide a theoretical framework to study optimality guarantees of sampling-based algorithms for the DMP problem by focusing on *driftless* control-affine dynamical systems of the form $\dot{\mathbf{x}}(t) = \sum_{i=1}^m g_i(\mathbf{x}(t))\mathbf{u}_i(t)$. While this model is representative for a large class of robotic systems (e.g., mobile robots with wheels that roll without slipping and multi-fingered robotic hands), it is of limited applicability in problems where momentum (i.e., drift) is a key feature of the problem setup (e.g., for a spacecraft or a helicopter).

Statement of Contributions: The objective of this paper is to provide a theoretical framework to study optimality guarantees of sampling-based algorithms for the DMP problem with *drift*. Specifically, as in [9], we focus on linear affine systems of the form

$$\dot{\mathbf{x}}[t] = A\mathbf{x}[t] + B\mathbf{u}[t] + \mathbf{c}, \quad \mathbf{x}[t] \in \mathcal{M}, \mathbf{u}[t] \in \mathcal{U},$$

where \mathcal{M} and \mathcal{U} are the configuration and control spaces, respectively, and it is of interest to find an obstacle-free trajectory π that minimizes the mixed time/energy criterion

$$c[\pi] = \int_0^T (1 + \mathbf{u}[t]^T R \mathbf{u}[t]) dt,$$

where R is a positive definite matrix that weights control energy expenditure versus traversal time. Henceforth, we will refer to a DMP problem involving linear affine dynamics and a mixed time/energy cost criterion as Linear Quadratic DMP (LQDMP). The LQDMP problem is relevant to applications for two main reasons: (i) it models the “essential” features of a number of robotic systems (e.g., spacecraft in deep space, helicopters, or even ground vehicles), and (ii) its theoretical study forms the backbone for sampling-based approaches that rely on linearization of more complex underlying dynamics. From a theoretical and algorithmic standpoint, the LQDMP problem presents two challenging features: (i) dynamics are not symmetric [1], which makes forward and backward reachable sets different and requires a more sophisticated analysis of sampling volumes to prove asymptotic optimality, and (ii) not all directions of motion are equivalent, in the sense that some motions incur dramatically higher cost than others due to the algebraic structure of the constraints. Indeed, these are the very same challenges that make the DMP problem with drift difficult in the first place, and they make approximation arguments (e.g., those needed to prove asymptotic optimality) more involved. Fortunately, for LQDMP an explicit characterization for the optimal trajectory connecting two sampled points in the absence of obstacles is available, which provides a foothold to begin the analysis. Specifically, the contribution of this paper is threefold. First, we show that *any* trajectory in an LQDMP problem may be “traced” arbitrarily well, with high probability, by connecting randomly distributed points from a sufficiently large sample set covering the configuration space. We will refer to this property as *probabilistic exhaustivity*, as opposed to probabilistic completeness [1], where the requirement is that *at least* one trajectory is traced with a sufficiently large sample set. Second, we introduce a sampling-

based algorithm for solving the LQDMP problem, namely the Differential Fast Marching Tree algorithm (DFMT*), whose design is enabled by our analysis of the notion of probabilistic exhaustivity. In particular, we are able to give a precise characterization of neighborhood radius, an important parameter for many asymptotically optimal motion planners, in contrast with previous work on LQDMP [9]. Third, by leveraging probabilistic exhaustivity, we show that DFMT* is asymptotically optimal. This analysis framework builds upon [10], and elements of our approach are inspired by [9]. We note that in [9], the authors present an excellent extension of RRT* that successfully solves the LQDMP problem in simulations, even when extended to linearized systems. The asymptotic optimality claim, however, relies only on a near-neighbor set size argument: we aim to put the analysis of the LQDMP problem on more rigorous theoretical footing.

Organization: This paper is structured as follows. In Section II we formally define the DMP problem we wish to solve. In Section III we review known results about the problem of optimally connecting fixed initial and terminal states under linear affine dynamics with a quadratic cost function. Furthermore, we provide a simple, yet novel (to the best of our knowledge) asymptotic characterization of the spectrum of the weighted controllability Gramian, which is instrumental to our analysis. In Section IV we prove the aforementioned probabilistic exhaustivity property for drift systems with linear affine dynamics. In Section V we present the DFMT* algorithm, and in Section VI we discuss its asymptotic optimality (together with a convergence rate characterization). Section VII contains proof-of-concept simulations. Finally, in Section VIII we discuss several features of our analysis, we draw some conclusions, and we discuss directions for future work.

II. PROBLEM FORMULATION

Let $\mathcal{M} \subseteq \mathbb{R}^n$ and $\mathcal{U} \subseteq \mathbb{R}^m$ be the configuration space and control space, respectively, of a robotic system. Within this space let us assume the dynamics of the robot are given by the linear affine system

$$\dot{\mathbf{x}}[t] = A\mathbf{x}[t] + B\mathbf{u}[t] + \mathbf{c}, \quad \mathbf{x}[t] \in \mathcal{M}, \mathbf{u}[t] \in \mathcal{U}, \quad (1)$$

where $A \in \mathbb{R}^{n \times n}$, $B \in \mathbb{R}^{n \times m}$, and $\mathbf{c} \in \mathbb{R}^n$ are constants.

A tuple $\pi = (\mathbf{x}[\cdot], \mathbf{u}[\cdot], T)$ defines a *dynamically feasible* trajectory, alternatively path, if the state evolution $\mathbf{x} : [0, T] \rightarrow \mathcal{M}$ and control input $\mathbf{u} : [0, T] \rightarrow \mathcal{U}$ satisfy equation (1) for all $t \in [0, T]$. We define the cost of a trajectory π by the function

$$c[\pi] = \int_0^T (1 + \mathbf{u}[t]^T R \mathbf{u}[t]) dt \quad (2)$$

where $R \in \mathbb{R}^{m \times m}$ is symmetric positive definite, constant, and given. We may rewrite this cost function as $c[\pi] = T + c_u[\pi]$, where $c_u[\pi] = \int_0^T \mathbf{u}[t]^T R \mathbf{u}[t] dt$, with the interpretation that this cost function penalizes both trajectory duration T and control effort c_u . The matrix R determines the relative costs of the control inputs, as well as their costs relative to the duration of the trajectory. We denote this linear affine dynamical system with cost by $\Sigma = (A, B, \mathbf{c}, R)$.

Let $\mathcal{M}_{\text{obs}} \subset \mathcal{M}$ be the obstacle region within the configuration space and consider the closed obstacle-free space $\mathcal{M}_{\text{free}} = \text{cl}[\mathcal{M} \setminus \mathcal{M}_{\text{obs}}]$. The starting configuration \mathbf{x}_{init} is

an element of $\mathcal{M}_{\text{free}}$, and the goal region $\mathcal{M}_{\text{goal}}$ is an open subset of $\mathcal{M}_{\text{free}}$. The trajectory planning problem is denoted by the tuple $(\Sigma, \mathcal{M}_{\text{free}}, \mathbf{x}_{\text{init}}, \mathcal{M}_{\text{goal}})$. A dynamically feasible trajectory $\pi = (\mathbf{x}, \mathbf{u}, T)$ is *collision-free* if $\mathbf{x}[t] \in \mathcal{M}_{\text{free}}$ for all $t \in [0, T]$. A trajectory π is said to be *feasible* for the trajectory planning problem $(\Sigma, \mathcal{M}_{\text{free}}, \mathbf{x}_{\text{init}}, \mathcal{M}_{\text{goal}})$ if it is dynamically feasible, collision-free, $\mathbf{x}[0] = \mathbf{x}_{\text{init}}$, and $\mathbf{x}[T] \in \mathcal{M}_{\text{goal}}$.

Let Π be the set of all feasible paths. The objective is to find the feasible path with minimum associated cost. We define the optimal trajectory planning problem as follows:

LQDMP problem: Given a trajectory planning problem $(\Sigma, \mathcal{M}_{\text{free}}, \mathbf{x}_{\text{init}}, \mathcal{M}_{\text{goal}})$ with cost function $c : \Pi \rightarrow \mathbb{R}_{\geq 0}$ given by equation (2), find a feasible path π^* such that $c[\pi^*] = \inf\{c[\pi] \mid \pi \text{ is feasible}\}$.

If no such path exists, report failure.

Our analysis will rely on two key sets of assumptions, relating, respectively, to the system Σ and the problem-specific parameters $\mathcal{M}_{\text{free}}, \mathbf{x}_{\text{init}}, \mathcal{M}_{\text{goal}}$.

Assumptions on system: We assume that the system Σ is controllable, (i.e., the pair (A, B) is controllable) [12] so that even disregarding obstacles there exist dynamically feasible trajectories between states.¹ Also, we assume that the control space is unconstrained, i.e. $\mathcal{U} = \mathbb{R}^m$, and that the cost weight matrix R is symmetric positive definite, so that every control direction has positive cost. These assumptions will be collectively referred to as A_Σ .

Assumptions on problem parameters: We require that the configuration space is a compact subset of \mathbb{R}^n so that we may sample from it. Furthermore, we require that the goal region $\mathcal{M}_{\text{goal}}$ has *regular boundary*, that is there exists $\xi > 0$ such that for almost all $\mathbf{y} \in \partial\mathcal{M}_{\text{goal}}$, there exists $\mathbf{z} \in \mathcal{M}_{\text{goal}}$ with $B[\mathbf{z}, \xi] \subseteq \mathcal{M}_{\text{goal}}$ and $\mathbf{y} \in \partial B[\mathbf{z}, \xi]$, where B denotes the Euclidean 2-norm ball. This requirement that the boundary of the goal region has bounded curvature almost everywhere ensures that a sampling procedure may expect to select points in the goal region near any point on the region's boundary. We make requirements on the *clearance* of the optimal trajectory, i.e., its “distance” from \mathcal{M}_{obs} [10]. For a given $\delta > 0$, the δ -interior of $\mathcal{M}_{\text{free}}$ is the set of all states that are at least a Euclidean distance δ away from any point in \mathcal{M}_{obs} . A collision-free path π is said to have strong δ -clearance if its state trajectory \mathbf{x} lies entirely inside the δ -interior of $\mathcal{M}_{\text{free}}$. A collision-free path π is said to have weak δ -clearance if there exists a path π' that has strong δ -clearance and there exists a homotopy ψ , with $\psi[0] = \pi$ and $\psi[1] = \pi'$ that satisfies the following three properties: (a) $\psi[\alpha]$ is a dynamically feasible trajectory for all $\alpha \in (0, 1]$, (b) $\lim_{\alpha \rightarrow 0} c[\psi[\alpha]] = c[\pi]$, and (c) for all $\alpha \in (0, 1]$ there exists $\delta_\alpha > 0$ such that $\psi[\alpha]$ has strong δ_α -clearance. Properties (a) and (b) are required since pathological obstacle sets may be constructed that squeeze all optimum-approximating homotopies into undesirable motion. In practice, however, as long as $\mathcal{M}_{\text{free}}$ does not contain any passages of infinitesimal width, the fact that Σ is controllable will allow every trajectory to be weak δ -clear. We claim that these assumptions about the problem parameters are mild, and can be regarded as “minimum” regularity assumptions.

¹This system controllability assumption is why we do not fold the constant drift term \mathbf{c} into the state \mathbf{x} .

All trajectories discussed in this paper are dynamically feasible unless otherwise noted. The symbol $\|\cdot\|$ denotes the 2-norm, induced or otherwise. The asymptotic notations O, Ω, Θ, o mean bounded above, bounded below, bounded both above and below, and asymptotically dominated, respectively.

III. OPTIMAL CONTROL IN THE ABSENCE OF OBSTACLES

The goal of this section is twofold: to review results about two-point boundary value problems for linear affine systems, and to present a simple, yet novel asymptotic characterization of the spectrum of the controllability Gramian. Both results will be instrumental to our analysis of LQDMP.

A. Two Point Boundary Value Problem

The material in this section is standard, we provide it to make the paper self-contained. Our presentation follows the treatment in [13, 9]. Specifically, this section is concerned with local steering between states in the absence of environment boundaries and obstacles. Given a start state $\mathbf{x}_0 \in \mathcal{M}$ and an end state $\mathbf{x}_1 \in \mathcal{M}$, the *two point boundary value problem* (2BVP) is to find a trajectory $\pi = (\mathbf{x}, \mathbf{u}, \tau^*)$ between $\mathbf{x}[0] = \mathbf{x}_0$ and $\mathbf{x}[\tau^*] = \mathbf{x}_1$ that satisfies the system Σ and minimizes its cost function (2). Denote this trajectory and its cost as $\pi^*[\mathbf{x}_0, \mathbf{x}_1]$ and $c^*[\mathbf{x}_0, \mathbf{x}_1]$ respectively:

$$\begin{aligned} \pi^*[\mathbf{x}_0, \mathbf{x}_1] &= \operatorname{argmin}\{\pi \mid \mathbf{x}[0] = \mathbf{x}_0 \wedge \mathbf{x}[\tau^*] = \mathbf{x}_1\} c[\pi] \\ c^*[\mathbf{x}_0, \mathbf{x}_1] &= \min\{\pi \mid \mathbf{x}[0] = \mathbf{x}_0 \wedge \mathbf{x}[\tau^*] = \mathbf{x}_1\} c[\pi]. \end{aligned}$$

Let us define the weighted controllability Gramian $G[t]$ as the solution of the Lyapunov equation

$$\dot{G}[t] = AG[t] + G[t]A^T + BR^{-1}B^T, \quad G[0] = 0,$$

which has the closed form expression

$$G[t] = \int_0^t \exp[As]BR^{-1}B^T \exp[A^Ts]ds. \quad (3)$$

Under the assumptions A_Σ (in particular, system (1) is controllable), we have that $G[t]$ is symmetric positive definite for all $t > 0$. This fact allows us to define the weighted norm $\|\cdot\|_{G^{-1}}$ for $\mathbf{x} \in \mathbb{R}^n$:

$$\|\mathbf{x}\|_{G[t]^{-1}} = \sqrt{\mathbf{x}^T G[t]^{-1} \mathbf{x}}.$$

Let $\bar{\mathbf{x}}[t]$ be the zero input response of system (1), that is the solution of the differential equation

$$\dot{\bar{\mathbf{x}}}[t] = A\bar{\mathbf{x}}[t] + \mathbf{c}, \quad \bar{\mathbf{x}}[0] = \mathbf{x}_0,$$

which has the closed form expression

$$\bar{\mathbf{x}}[t] = \exp[At]\mathbf{x}_0 + \int_0^t \exp[As]\mathbf{c}ds. \quad (4)$$

Then for a fixed arrival time τ the optimal control policy for the fixed-time 2BVP is given by [13]:

$$\mathbf{u}[t] = R^{-1}B^T \exp[A^T(\tau - t)]G[\tau]^{-1}(\mathbf{x}_1 - \bar{\mathbf{x}}[\tau]), \quad (5)$$

which corresponds to the minimal cost (as a function of travel time τ)

$$c[\tau] = \tau + \|\mathbf{x}_1 - \bar{\mathbf{x}}[\tau]\|_{G[\tau]^{-1}}^2. \quad (6)$$

The optimal connection time τ^* may be computed by minimizing (6) over τ . The state trajectory $\mathbf{x}[t]$ that evolves from this control policy may be computed explicitly as:

$$\mathbf{x}[t] = \bar{\mathbf{x}}[t] + G[t] \exp[A^T(\tau^* - t)]G[\tau^*]^{-1}(\mathbf{x}_1 - \bar{\mathbf{x}}[\tau^*]). \quad (7)$$

Let $\pi^*[\mathbf{x}_0, \mathbf{x}_1, \dots, \mathbf{x}_K]$ denote the concatenation of the trajectories $\pi^*[\mathbf{x}_k, \mathbf{x}_{k+1}]$ between successive states $\mathbf{x}_0, \mathbf{x}_1, \dots, \mathbf{x}_K \in \mathcal{M}$.

B. Small-Time Characterization of the Spectrum of the Controllability Gramian

We begin by briefly reviewing the concept of *controllability indices*² for a controllable system (A, B) . Let \mathbf{b}_k denote the k th column of B . Consider searching the columns of the controllability matrix $\mathcal{C}[A, B] = [B \ AB \ \dots \ A^{n-1}B]$ from left to right for a set of n linearly independent vectors. This process is well-defined for a controllable pair (A, B) since $\text{rank}[\mathcal{C}[A, B]] = n$. The resulting set $\mathcal{S} = \{\mathbf{b}_1, A\mathbf{b}_1, \dots, A^{\nu_1-1}\mathbf{b}_1, \mathbf{b}_2, \dots, A^{\nu_2-1}\mathbf{b}_2, \dots, A^{\nu_m-1}\mathbf{b}_m\}$ defines the controllability indices $\{\nu_1, \dots, \nu_m\}$ where $\sum_{k=1}^m \nu_k = n$ and $\nu = \max_k \nu_k$ is called the *controllability index* of (A, B) . The ν_k give a fundamental notion of how difficult a system is to control in various directions; indeed these indices are a property of the system invariant with respect to similarity transformation, e.g. permuting the columns of B . We may also label the vectors of \mathcal{S} as $\mathbf{v}_1, \dots, \mathbf{v}_n$ in the order that they come up in $\mathcal{C}[A, B]$. That is, $\mathbf{v}_i = A^{e_i}\mathbf{b}_{k_i}$ and $e_i \leq e_j$ iff $i \leq j$ (note: $e_n = \nu - 1$). Let $\hat{\mathbf{v}}_1, \dots, \hat{\mathbf{v}}_n$ be an orthogonalization of the \mathbf{v}_i 's so that $V = \hat{V}X$ where V and \hat{V} have the \mathbf{v}_i 's and $\hat{\mathbf{v}}_i$'s as columns respectively, and X is upper triangular.

Lemma III.1 (Small-Time Gramian Asymptotics). *Let the eigenvalues of $G[t]$ be $\lambda_1[t] \geq \lambda_2[t] \geq \dots \geq \lambda_n[t]$. Then $\lambda_i[t] = \Theta[t^{2e_i+1}]$ as $t \rightarrow 0$ for $i = 1, \dots, n$.*

Proof. We apply the Courant-Fischer Theorem:

$$\lambda_i = \min_{\mathcal{V}} \{\max_{\mathbf{z} \in \mathcal{V}^*} \{R_{G[t]}[\mathbf{z}] \mid \dim[\mathcal{V}] = n - i + 1\}\} \quad (8)$$

where $R_{G[t]}[\mathbf{z}] = \mathbf{z}^T G[t] \mathbf{z} / \|\mathbf{z}\|$, \mathcal{V} denotes a linear subspace of \mathbb{R}^n , and $\mathcal{V}^* = \mathcal{V} \setminus \{\mathbf{0}\}$. Note that

$$\begin{aligned} \hat{\mathbf{v}}_i^T \exp[At]B &= \sum_{j=0}^{e_i-1} \frac{t^j}{j!} \hat{\mathbf{v}}_i^T A^j B + \sum_{j=e_i}^{\infty} \frac{t^j}{j!} \hat{\mathbf{v}}_i^T A^j B \\ &= \sum_{j=e_i}^{\infty} \frac{t^j}{j!} \hat{\mathbf{v}}_i^T A^j B = \Theta[t^{e_i}] \quad \text{as } t \rightarrow 0, \end{aligned}$$

because $\hat{\mathbf{v}}_i^T A^j B = 0$ for all $j < e_i$, $k = 1, \dots, m$ by construction. Then

$$\begin{aligned} R_{G[t]}[\hat{\mathbf{v}}_i] &= \int_0^t \hat{\mathbf{v}}_i^T \exp[As] B R^{-1} B^T \exp[A^T s] \hat{\mathbf{v}}_i \, ds \\ &= \int_0^t \Theta[s^{2e_i}] \, ds = \Theta[t^{2e_i+1}] \quad \text{as } t \rightarrow 0. \end{aligned}$$

Making the identification $\mathcal{V} = \text{span}\{\hat{\mathbf{v}}_i, \dots, \hat{\mathbf{v}}_n\}$ in Equation (8) implies that $\lambda_i = O[t^{2e_i+1}]$; to see that $\lambda_i = \Omega[t^{2e_i+1}]$ we note that any subspace \mathcal{V} of dimension $n - i + 1$ cannot satisfy $\hat{\mathbf{v}}_1, \dots, \hat{\mathbf{v}}_i \in \mathcal{V}^\perp$, as $\dim[\mathcal{V}^\perp] = i - 1$. \square

Lemma III.1 has three immediate corollaries. The first upper bounds $\|G[t]^{-1}\|$ which bounds the local cost of motion in any direction. The second relates $\|\cdot\|_{G[t]^{-1}}$ to the Euclidean norm $\|\cdot\|$ through a norm-equivalence inequality. The third is a lower bound for the determinant of $G[t]$, a result that will prove useful for estimating the volumes of reachable sets.

Lemma III.2 (Small-time minimum eigenvalue of controllability Gramian). *Suppose that the pair (A, B) has controllability index ν , then $\lambda_n[G[t]] = \Theta[t^{2\nu-1}]$ as $t \rightarrow 0$, or equivalently, $\|G[t]^{-1}\| = \Theta[t^{-2\nu+1}]$.*

²See [14, p. 431] or [12, p. 150] for a more detailed treatment.

Lemma III.3 (Norm Equivalence). *Suppose that the pair (A, B) has controllability index ν , and consider the Cholesky factorization $G[t] = L[t]L[t]^T$. Then for $\mathbf{x} \in \mathbb{R}^n$, $\|\mathbf{x}\|_{G[t]^{-1}} = \|L[t]^{-1}\mathbf{x}\|$ and*

$$\frac{\|\mathbf{x}\|}{\|L[t]\|} \leq \|\mathbf{x}\|_{G[t]^{-1}} \leq \|L[t]^{-1}\| \|\mathbf{x}\|$$

where $\|L[t]\| = \Theta[t^{-1/2}]$ and $\|L[t]^{-1}\| = \Theta[t^{-\nu+1/2}]$.

Lemma III.4 (Small-time determinant of controllability Gramian). *Suppose that the pair (A, B) has controllability indices $\{\nu_1, \dots, \nu_m\}$, then $\det[G[t]] = \Theta[t^D]$ as $t \rightarrow 0$ where $D = \sum_{k=1}^m \nu_k^2$.*

IV. PROBABILISTIC EXHAUSTIVITY

In this section we prove a key result characterizing random sampling schemes for the LQDMP problem: *any* feasible trajectory through the configuration space \mathcal{M} is “traced” arbitrarily well by connecting randomly distributed points from a sufficiently large sample set covering the configuration space. We will refer to this property as probabilistic exhaustivity. The same notion of probabilistic exhaustivity (clearly much stronger than the usual notion of probabilistic completeness) was introduced in the related paper [10] in the context of DMP for *driftless* systems. The result proven in that work does not carry over to the drift case as it relies on the metric inequality to bound the cost of approximate paths; the drift case lacks the control symmetry to make such estimates. Thus in order to prove probabilistic exhaustivity in the case of linear affine systems, we first provide a result analogous to the metric inequality characterizing the effect that perturbations of the endpoints of a path have on its cost and state trajectory. The idea, then, is that tracing waypoints may be selected as small perturbations of points along the trajectory to be approximated, provided the sample density is high enough.

Lemma IV.1 (Fixed-Time Local Trajectory Approximation). *Let $\mathbf{x}_0, \mathbf{x}_1 \in \mathcal{M}$, $\mathbf{x}_0 \neq \mathbf{x}_1$, $\pi = \pi^*[\mathbf{x}_0, \mathbf{x}_1] = (\mathbf{x}, \mathbf{u}, \tau^*)$, and denote $c = c[\pi]$. Consider bounded start and end state perturbations $\delta\mathbf{x}_0, \delta\mathbf{x}_1 \in \mathbb{R}^n$ such that $\max\{\|\delta\mathbf{x}_0\|_{G[\tau^*]^{-1}}, \|\delta\mathbf{x}_1\|_{G[\tau^*]^{-1}}\} \leq \eta\sqrt{c}$. Let $\sigma = \pi^*[\mathbf{x}_0 + \delta\mathbf{x}_0, \mathbf{x}_1 + \delta\mathbf{x}_1] = (\mathbf{y}, \mathbf{v}, \tilde{\tau}^*)$ be the optimal trajectory between the perturbed endpoints. Then for $\mathbf{x}_0, \mathbf{x}_1$ such that c is sufficiently small, we have the cost bound*

$$c[\sigma] \leq c[\pi] (1 + 4\eta + O[\eta^2 + \eta c]).$$

Additionally we may bound the geometric extent of σ :

$$\|\mathbf{y}[t] - \mathbf{x}_0\| = O[c(\|\mathbf{x}_0\| + \eta + 1)]$$

for $t \in [0, \tilde{\tau}^*]$.

Proof. Since σ is the optimum, regardless of the value of $\tilde{\tau}^*$ we have the upper bound

$$c[\sigma] \leq \tau^* + \|\mathbf{x}_1 - \bar{\mathbf{x}}[\tau^*] + \delta\mathbf{x}_1 - \exp[A\tau^*]\delta\mathbf{x}_0\|_{G[\tau^*]^{-1}}^2$$

which expands as

$$\begin{aligned} c[\sigma] - c &\leq 2(\mathbf{x}_1 - \bar{\mathbf{x}}[\tau^*])^T G[\tau^*]^{-1} (\delta\mathbf{x}_1 - \exp[A\tau^*]\delta\mathbf{x}_0) \\ &\quad + \|\delta\mathbf{x}_1 - \exp[A\tau^*]\delta\mathbf{x}_0\|_{G[\tau^*]^{-1}}^2 \\ &\leq 2\|\mathbf{x}_1 - \bar{\mathbf{x}}[\tau^*]\|_{G[\tau^*]^{-1}} \|\delta\mathbf{x}_1 - \exp[A\tau^*]\delta\mathbf{x}_0\|_{G[\tau^*]^{-1}} \\ &\quad + O[\eta^2 c] \\ &\leq 2\sqrt{c} - \tau^* (\eta\sqrt{c} + \eta\sqrt{c}(1 + O[\tau^*])) + O[\eta^2 c] \\ &\leq 4\eta c(1 + O[\eta + c]), \end{aligned}$$

where in the second line we have applied the (weighted) Cauchy-Schwarz inequality, and, in the last line, the fact that $\tau^* \leq c$.

To bound $\|\mathbf{y}[t] - \mathbf{x}_0\|$, we first apply Cauchy-Schwarz:

$$\begin{aligned} \int_0^{\tau^*} \|\mathbf{v}[t]\| dt &\leq \left(\tilde{\tau}^* \int_0^{\tilde{\tau}^*} \|\mathbf{v}[t]\|^2 dt \right)^{1/2} \\ &\leq \left(\frac{\tilde{\tau}^*}{\lambda_{\min}(R)} \int_0^{\tilde{\tau}^*} \mathbf{v}[t]^T R \mathbf{v}[t] dt \right)^{1/2} \\ &= \left(\frac{\tilde{\tau}^*}{\lambda_{\min}(R)} (c[\sigma] - \tilde{\tau}^*) \right)^{1/2} \leq \frac{c[\sigma]}{2\sqrt{\lambda_{\min}(R)}}. \end{aligned}$$

Then integrating the system dynamics (1) yields

$$\begin{aligned} \|\mathbf{y}[t] - \mathbf{x}_0\| &= \left\| (e^{At} - I)\mathbf{x}_0 + e^{At}\delta\mathbf{x}_0 + \int_0^t e^{As}\mathbf{c} ds \right. \\ &\quad \left. + \int_0^t e^{A(t-s)}B\mathbf{v}(s) ds \right\| \\ &= O \left[\tilde{\tau}^* \|\mathbf{x}_0\| + \tau^{1/2} \|\delta\mathbf{x}_0\|_{G[\tau^*]^{-1}} + \tilde{\tau}^* + \int_0^{\tilde{\tau}^*} \|\mathbf{v}[t]\| dt \right] \\ &= O[c(\|\mathbf{x}_0\| + \eta + 1)], \end{aligned}$$

making use of a norm equivalence bound (Lemma III.3) in line two. We note that the asymptotic constants depend only on the fixed system dynamics $\Sigma = (A, B, \mathbf{c}, R)$. \square

Motivated by Lemma IV.1, we define the perturbation ball

$$\Delta[\mathbf{x}, \tau, r] = \left\{ \mathbf{z} \mid \|\mathbf{x} - \mathbf{z}\|_{G[\tau]^{-1}} \leq r \right\}.$$

This set represents perturbations of \mathbf{x} with limited effects on both incoming and outgoing trajectories (depending on whether a point is viewed as an end state or start state perturbation respectively). We note that since $\|\cdot\|_{G[t]^{-1}}$ is decreasing as t increases, we have

$$\Delta[\mathbf{x}, \tau_1, r] \subset \Delta[\mathbf{x}, \tau_2, r] \quad \text{for } \tau_1 \leq \tau_2. \quad (9)$$

To understand how often sample points of a planning algorithm will lie within $\Delta[\mathbf{x}, \tau, r]$, we lower bound its volume.

Remark IV.2 (Bounding Perturbation Ball Volume). *The inequality $\|\mathbf{x} - \mathbf{z}\|_{G[\tau]^{-1}} \leq r$ defines an ellipse with volume*

$$\mu[\Delta[\mathbf{x}, \tau, r]] = r^n \zeta_n \sqrt{\det[G[\tau]]}$$

where ζ_n denotes the volume of the unit ball in \mathbb{R}^n . Given our asymptotic characterization of $\det[G[\tau]]$ in Lemma III.4, there is a threshold $\tau_\mu > 0$ and constant $C_\mu > 0$ such that

$$\mu[\Delta[\mathbf{x}, \tau, r]] \geq C_\mu r^n \tau^{D/2}$$

for all $\tau < \tau_\mu$.

To ensure that the $\tau^{D/2}$ term above does not become vanishingly small in application with Lemma IV.1, we also lower bound connection time in terms of connection cost.

Lemma IV.3 (Optimal Cost/Time Breakdown). *Let $\mathbf{x}_0, \mathbf{x}_1 \in \mathcal{M}$, $\mathbf{x}_0 \neq \mathbf{x}_1$, $\pi = \pi^*[\mathbf{x}_0, \mathbf{x}_1] = (\mathbf{x}, \mathbf{u}, \tau^*)$, and denote $c = c[\pi]$. For $\mathbf{x}_0, \mathbf{x}_1$ such that c is sufficiently small, $\tau^* \geq \frac{c}{3}$.*

Proof. For fixed $\mathbf{x}_0 \neq \mathbf{x}_1$, consider the series expansion of the control effort term $\|\mathbf{x}_1 - \bar{\mathbf{x}}[\tau]\|_{G[\tau]^{-1}}^2$ about $\tau = 0$. We claim that $\|\mathbf{x}_1 - \bar{\mathbf{x}}[\tau]\|_{G[\tau]^{-1}}^2 = a\tau^{-e} + O[\tau^{-e+1}]$ for some $1 \leq e \leq 2\nu - 1$. The fact that $e \geq 1$ follows from the fact that $\mathbf{x}_1 - \bar{\mathbf{x}}[\tau] = (\mathbf{x}_1 - \mathbf{x}_0) + O[\tau]$ has a nonzero zeroth order term and $\lambda_{\min}[G[\tau]^{-1}] = \Theta[\tau^{-1}]$ as a consequence of Lemma III.1. We note that for general $\mathbf{x}_0, \mathbf{x}_1$, however,

it is almost certain that $(\mathbf{x}_1 - \mathbf{x}_0)$ or one of the low-order (in τ) terms of $\mathbf{x}_1 - \bar{\mathbf{x}}[\tau]$ will have a component along the maximal eigenvector of $G[\tau]^{-1}$, which may result in a series expansion term with e up to $2\nu - 1$. Then

$$c[\tau] = \tau + a\tau^{-e} + O[\tau^{-e+1}]$$

is maximized at $\tau^* = (ae)^{1/(1+e)} + o(1)$ and we compute the ratio

$$\tau^*/c = (1 + 1/e)^{-1}(1 + o(1))$$

as $c \rightarrow 0$. The dominant term in the asymptotics is smallest when $e = 1$ (corresponding to $\tau^*/c \approx 1/2$); in particular we have $\tau^*/c \geq 1/3$ for $\mathbf{x}_0, \mathbf{x}_1$ with c sufficiently small. \square

Lemma IV.1 is a statement about local trajectory approximation. We now define what it means for a series of states to closely approximate a given global trajectory. Let $\pi = (\mathbf{x}, \mathbf{u}, T_\pi)$ be a dynamically feasible trajectory. Given a set of waypoints $\{\mathbf{y}_k\}_{k=0}^K \subset \mathcal{M}$, we associate the trajectory $\sigma = \pi^*[\mathbf{y}_0, \dots, \mathbf{y}_K] = (\mathbf{y}, \mathbf{v}, T_\sigma)$. We consider the $\{\mathbf{y}_k\}$ to (ε, r, p) -trace the trajectory π if: (a) the cost of σ is bounded as $c[\sigma] \leq (1 + \varepsilon)c[\pi]$, (b) $c^*[\mathbf{y}_k, \mathbf{y}_{k+1}] \leq r$ for all k , and (c) the maximum distance from any point of \mathbf{y} to \mathbf{x} is no more than p , i.e. $\max_{t \in [0, T_\sigma]} (\min_{s \in [0, T_\pi]} \|\mathbf{y}[t] - \mathbf{x}[s]\|) \leq p$. The combination of these three properties is what makes σ , if approximating a near-globally-optimal trajectory π , amenable to recovery by the path planning algorithms we propose in the next section. In particular, (b) ensures that σ is the concatenation of uniformly local connections. In Theorem IV.6 we show that suitable waypoints may be found with high probability as a subset of a set of randomly sampled nodes, the proof of which requires the following two technical lemmas lower bounding the probability that a sample set will provide adequate coverage around a trajectory of interest. Let $\text{SampleFree}[N]$ denote a set of N points sampled independently and identically from the uniform distribution on $\mathcal{M}_{\text{free}}$.

Lemma IV.4 (Lemma IV.3, [10]). *Fix $N \in \mathbb{N}$, $\alpha \in (0, 1)$, and let S_0, \dots, S_K be disjoint subsets of $\mathcal{M}_{\text{free}}$ with*

$$\mu[S_k] = \mu[S_1] \geq (2 + \log(1/\alpha))e^2 \left(\frac{1}{N} \right) \mu[\mathcal{M}_{\text{free}}],$$

for each k . Let $V = \text{SampleFree}[N]$; then the probability that more than an α fraction of the sets S_k contain no point of V is bounded as:

$$\mathbb{P}[\#\{k \in \{0, \dots, K\} : S_k \cap V = \emptyset\} \geq \alpha K] \leq 2e^{-\alpha K}.$$

Lemma IV.5 (Lemma IV.4, [10]). *Fix $N \in \mathbb{N}$ and let T_0, \dots, T_K be subsets of $\mathcal{M}_{\text{free}}$, possibly overlapping, with*

$$\mu[T_k] = \mu[T_1] \geq \kappa \left(\frac{\log N}{N} \right) \mu[\mathcal{M}_{\text{free}}]$$

for each k and some constant $\kappa > 0$. Let $V = \text{SampleFree}[N]$; then the probability that there exists a T_k that does not contain a point of V is bounded as:

$$\mathbb{P} \left[\bigvee_{k=0}^K \{T_k \cap V = \emptyset\} \right] \leq KN^{-\kappa}.$$

The proofs of these two lemmas may be found in our related work [10]. As in that work and [15], our approach here for proving probabilistic exhaustivity proceeds by tiling

³We note that this bound does not depend on the actual values of \mathbf{x}_0 and \mathbf{x}_1 (in particular the constant a), but only their optimal connection cost.

the span of a path to be traced with two sequences of concentric perturbation balls: a sequence of “small” balls and a sequence of “large” balls. With high probability, all but a tiny α fraction of the small balls will contain a point from the sample set (Lemma IV.4), and for any small balls that do not contain such a point we ensure that the corresponding large ball does (Lemma IV.5). We take these points as a sequence of waypoints which tightly follows the reference path with few exceptions, and never has a gap over any section of the reference path when it does deviate further.

Theorem IV.6 (Probabilistic exhaustivity). *Let Σ be a system satisfying the assumptions A_Σ and suppose $\pi = (\mathbf{x}, \mathbf{u}, T)$ is a dynamically feasible trajectory with strong δ -clearance, $\delta > 0$. Let $N \in \mathbb{N}$, $\varepsilon > 0$, and consider a set of sample nodes $V = \{\mathbf{x}[0]\} \cup \text{SampleFree}[N]$. Define $\tilde{D} = (n + D)/2$, $C_{\Sigma, \mathcal{M}_{\text{free}}} = \left(C_\mu^{-1} \tilde{D}^{-1} 6^{n+D/2} 2^{n/2} \mu[\mathcal{M}_{\text{free}}]\right)^{1/\tilde{D}}$ and consider the event E_N that there exist waypoints $\{\mathbf{y}_k\}_{k=0}^K \subset V$ which (ε, r_N, p_N) -trace π , where*

$$r_N = (1 + \eta)^{1/\tilde{D}} C_{\Sigma, \mathcal{M}_{\text{free}}} \left(\frac{\log N}{N} \right)^{1/\tilde{D}},$$

for a free parameter $\eta \geq 0$, and $p_N = C_p r_N$ for some constant C_p . Then, as $N \rightarrow \infty$, the probability that no such waypoint set exists is asymptotically bounded as

$$1 - \mathbb{P}[E_N] = O\left(N^{-\eta/\tilde{D}} \log^{-1/\tilde{D}} N\right).$$

Proof. Note that in the case $c[\pi] = 0$ we may pick $\mathbf{y}_0 = \mathbf{x}[0]$ to be the only waypoint and the result is trivial. Therefore assume $c[\pi] > 0$. Make the identification $\alpha = \varepsilon/4$, $\beta = \varepsilon/2$ and fix N sufficiently large so that $r_N/6 \leq \tau_\mu$ and also:

$$\log N \geq \beta^{-n} \tilde{D} (2 - \log(\alpha)) e^2 / (1 + \eta). \quad (10)$$

Take $\mathbf{x}[t_k]$ to be points spaced along π at cost intervals $r_N/2$; more precisely let $t_0 = 0$, and for $k = 1, 2, \dots$ consider

$$t_k = \min \{t \in (t_k, T) \mid c^*[\mathbf{x}[t_{k-1}], \mathbf{x}[t]] \geq r_N/2\}.$$

Let K be the first k for which the set is empty; take $t_K = T$. Note that by construction, we have $K \leq \lceil 2c[\pi]/r_N \rceil$.

We consider the sets $T_k = \Delta[\mathbf{x}[t_k], r_N/6, (1/6)\sqrt{r_N/2}]$ and $S_k = \Delta[\mathbf{x}[t_k], r_N/6, (\beta/6)\sqrt{r_N/2}]$. In particular the time $r_N/6$ here is chosen so that, by Lemma IV.3, the optimal connection times between the \mathbf{x}_k satisfy $\tau^*[\mathbf{x}[t_{k-1}], \mathbf{x}[t_k]] \geq c^*[\mathbf{x}[t_{k-1}], \mathbf{x}[t_k]]/3 = r_N/6$. Applying the ball containment property (9) this means that for any such τ^* , $\Delta[\mathbf{x}[t_k], r_N/6, \rho] \subset \Delta[\mathbf{x}[t_k], \tau^*, \rho]$ for $\rho = \sqrt{r_N/2}$ or $(\beta/4)\sqrt{r_N/2}$. From Remark IV.2 and our choice of r_N we have the volume bound

$$\mu(T_k) \geq \mu[\mathcal{M}_{\text{free}}] \left(\frac{1 + \eta}{\tilde{D}} \right) \left(\frac{\log N}{N} \right), \quad (11)$$

and similarly

$$\mu(S_k) \geq \beta^n \mu[\mathcal{M}_{\text{free}}] \left(\frac{1 + \eta}{\tilde{D}} \right) \left(\frac{\log N}{N} \right), \quad (12)$$

for each k . Combining equation (11) and Lemma IV.5, we have that the probability that there exists a T_k that does not contain a sample point (i.e. $T_k \cap V = \emptyset$) is bounded as:

$$\mathbb{P}\left[\bigvee_{m=0}^K \{T_m \cap V = \emptyset\}\right] \leq K N^{-(1+\eta)/\tilde{D}}.$$

We note that the S_k are disjoint (as long as $\varepsilon < 2\sqrt{6}$) since $c^*[\mathbf{x}[t_{k-1}], \mathbf{x}[t_k]] = r_N/2$ implies

$\|\mathbf{x}[t_{k-1}] - \mathbf{x}[t_k]\|_{G[r_N/6]^{-1}} \geq \sqrt{r_N/3}$. Then we may combine equations (10) and (12), which together imply that the S_k satisfy the condition of Lemma IV.4, to see that the probability that more than an α fraction of the S_k do not contain a sample point is bounded as:

$$\mathbb{P}[\#\{k \in \{0, \dots, K\} : S_k \cap V = \emptyset\} \geq \alpha K] \leq 2e^{-\alpha K}.$$

Now, as long as neither of these possibilities holds (i.e. if every T_k and at least a $(1 - \alpha)$ fraction of the S_k contains a point of V), we will show that the existence of suitable waypoints $\{\mathbf{y}_k\}_{k=0}^K \subset V$ is guaranteed. In that case then we may union bound the probability of failure:

$$1 - \mathbb{P}[E_N] \leq K N^{-(1+\eta)/\tilde{D}} + 2e^{-\alpha K} \\ = O\left[N^{-\eta/\tilde{D}} \log^{-1/\tilde{D}} N\right],$$

as $N \rightarrow \infty$ (the first term dominates asymptotically), where we have used the fact that $K \leq \lceil 2c[\pi]/r_N \rceil = O[(N/\log N)^{1/\tilde{D}}]$.

Suppose that every T_k and at least a $(1 - \alpha)$ fraction of the S_k contains a point of V . Choose points $\{\mathbf{y}_k\} \subset V$ accordingly: within S_k if possible, and within T_k otherwise. We may apply Lemma IV.1 to verify that these points (ε, r_N, p_N) -trace \mathbf{x} . For $\mathbf{y}_k \in S_k, \mathbf{y}_{k+1} \in S_{k+1}$ we have:

$$c^*[\mathbf{y}_{k-1}, \mathbf{y}_k] \leq (1 + \beta) c^*[\mathbf{x}[t_{k-1}], \mathbf{x}[t_k]] \\ = (1 + \varepsilon/2)(r_N/2).$$

Since all but a $2\alpha = \varepsilon/2$ fraction of successive points $\mathbf{y}_{k-1}, \mathbf{y}_k$ must both be in S sets and obey the above cost bound, and the remaining pairs satisfy the analogous bound for T sets (with 2 instead of $(1 + \varepsilon/2)$), the total cost of $\pi^*[\mathbf{y}_0, \dots, \mathbf{y}_K] = (\mathbf{y}, \mathbf{v}, T_\sigma)$ is bounded above by $(1 + \varepsilon)c[\pi]$. We also have $c^*[\mathbf{y}_{k-1}, \mathbf{y}_k] \leq r_N$ for all k . The maximum Euclidean distance from any point of \mathbf{y} (say, on the segment $\pi^*[\mathbf{y}_k, \mathbf{y}_{k+1}]$) to \mathbf{x} is bounded above by its distance to $\mathbf{x}[t_k]$, which by Lemma IV.1 is $O[r_N(\|\mathbf{x}_0\| + 1)] = O[r_N]$ as $N \rightarrow \infty$ since $\|\mathbf{x}[t]\|$ achieves some fixed maximum over $[0, T]$. \square

V. DFMT* ALGORITHM

The algorithm presented here is based on FMT*, from the recent work of [15], which can be thought of as an accelerated version of PRM* [4]. Briefly, PRM* first samples all the vertices, then constructs a fully *locally* connected graph, and then performs shortest path search (e.g., Dijkstra’s algorithm) on the graph to obtain a solution. FMT* also samples all vertices first, but instead of a graph, lazily builds a tree *via dynamic programming* that very closely approximates the shortest-path tree for PRM*, but saves a multiplicative factor of $O[\log(n)]$ collision-checks by not constructing the full graph. The algorithm given by Algorithm 1, DFMT*, is not fundamentally different from the original FMT* algorithm, but mainly changes what “local” means under differential constraints (similar to [10], but now with drift). One more difference of DFMT* presented here, even from the algorithm in [10], is that the edges are now directed, reflecting the fundamental asymmetry of differential constraints with drift.

Specifically, define the fixed-time forward-reachable and backwards-reachable sets respectively:

$$R^+[\mathbf{x}, r] = \{\mathbf{x}' \in \mathcal{M} \mid c^*[\mathbf{x}, \mathbf{x}'] < r\}$$

$$R^-[\mathbf{x}, r] = \{\mathbf{x}' \in \mathcal{M} \mid c^*[\mathbf{x}', \mathbf{x}] < r\}.$$

Membership in either reachable set may be checked by minimizing the explicit cost function (6) over travel time. The set of samples to check for membership may be pruned by considering the form of $G[t]$, as suggested in [9]. Let $\text{CollisionFree}[\mathbf{x}_1, \mathbf{x}_2]$ denote the boolean function which returns true if and only if $\pi^*[\mathbf{x}_1, \mathbf{x}_2]$ lies within $\mathcal{M}_{\text{free}}$. Given a set of vertices V , a state $\mathbf{x} \in \mathcal{M}$, and a cost threshold $r > 0$, let $\text{Near}^\pm[V, \mathbf{x}, r] = V \cap R^\pm[\mathbf{x}, r]$. Let $(\mathbf{x}_1, \mathbf{x}_2)$ denote the directed edge corresponding to $\pi^*[\mathbf{x}_1, \mathbf{x}_2]$ with edge weight $c^*[\mathbf{x}_1, \mathbf{x}_2]$. Given a directed graph $G = (V, E)$, where V is the vertex set and E is the edge set, and a vertex $\mathbf{x} \in V$, let $\text{Cost}[\mathbf{x}, G]$ be the function that returns the cost of the shortest (directed) path in the graph G between the vertices \mathbf{x}_{init} and \mathbf{x} . Let $\text{Path}[\mathbf{x}, G]$ be the function that returns the path achieving that cost. The DFMT* algorithm is given in Algorithm 1. The algorithm uses two mutually exclusive sets, namely H and W . The *unexplored* set W stores all samples in the sample set V that have not yet been considered for addition to the tree of paths. The *wavefront* set H , on the other hand, tracks in sorted order (by cost from the root) only those nodes which have already been added to the tree that are near enough to tree leaves to actually form better connections. A detailed description of the algorithm would parallel the one provided in [15] and is omitted due to space limitations, we refer the interested reader to [15]. An extension of PRM*, which we denote by DPRM*, may also be defined in a straightforward manner as in [10], although we omit the full description here. Briefly, DPRM* searches the graph of all local collision-free connections that appear in any *Near* set (as opposed to the tree subgraph constructed by DFMT*) for the least cost trajectory.

Algorithm 1 Differential Fast Marching Tree (DFMT*)

```

1  $V \leftarrow \{\mathbf{x}_{\text{init}}\} \cup \text{SampleFree}[N]; E \leftarrow \emptyset$ 
2  $W \leftarrow V \setminus \{\mathbf{x}_{\text{init}}\}; H \leftarrow \{\mathbf{x}_{\text{init}}\}; \mathbf{z} \leftarrow \mathbf{x}_{\text{init}}$ 
3 while  $\mathbf{z} \notin \mathcal{M}_{\text{goal}}$  do
4    $H_{\text{new}} \leftarrow \emptyset$ 
5    $X_{\text{near}} = \text{Near}^+[V \setminus \{\mathbf{z}\}, \mathbf{z}, r_N] \cap W$ 
6   for  $\mathbf{x} \in X_{\text{near}}$  do
7      $Y_{\text{near}} \leftarrow \text{Near}^-[V \setminus \{\mathbf{x}\}, \mathbf{x}, r_N] \cap H$ 
8      $\mathbf{y}_{\text{min}} \leftarrow \arg \min_{\mathbf{y} \in Y_{\text{near}}} \{\text{Cost}[\mathbf{y}, T = (V, E)] + c^*[\mathbf{y}, \mathbf{x}]\}$ 
9     if  $\text{CollisionFree}[\mathbf{y}_{\text{min}}, \mathbf{x}]$  then
10        $E \leftarrow E \cup \{(\mathbf{y}_{\text{min}}, \mathbf{x})\}$ 
11        $H_{\text{new}} \leftarrow H_{\text{new}} \cup \{\mathbf{x}\}$ 
12        $W \leftarrow W \setminus \{\mathbf{x}\}$ 
13     end if
14   end for
15    $H \leftarrow (H \cup H_{\text{new}}) \setminus \{\mathbf{z}\}$ 
16   if  $H = \emptyset$  then
17     return Algorithm Failure
18   end if
19    $\mathbf{z} \leftarrow \arg \min_{\mathbf{y} \in H} \text{Cost}[\mathbf{y}, T = (V, E)]$ 
20 end while
21 return  $\text{Path}[\mathbf{z}, T = (V, E)]$ 
```

VI. ASYMPTOTIC OPTIMALITY OF DFMT*

In this section, we state the asymptotic optimality of DFMT*, for which the asymptotic optimality of DPRM*

is a corollary. We note that in contrast to the work required to establish probabilistic exhaustivity for this class of differentially constrained systems and cost functions, the argument that DFMT* recovers paths at least as good as any waypoint-traced trajectory is essentially equivalent to the proofs presented in [10] and [15]. That is, the idea that DFMT* (or DPRM*) can connect closely spaced sample points at a resolution sufficiently fine that every connection takes place away from the influence of the obstacle set is not a feature specific to the LQDMP problem. Thus we state the following theorem and, in the interest of brevity, refer the reader to the proofs of Theorems VI.1 and VI.2 presented in [10]. We note that the following optimality result for DFMT* also provides a *convergence rate* bound, but to avoid confusion we emphasize that this bound is given in terms of sample size N . For a discussion of how sample size relates to run time for FMT*-style algorithms see [15].

Theorem VI.1 (DFMT* asymptotic optimality). *Let $(\Sigma, \mathcal{M}_{\text{free}}, \mathbf{x}_{\text{init}}, \mathcal{M}_{\text{goal}})$ be a trajectory planning problem satisfying the assumptions A_Σ and with $\mathcal{M}_{\text{goal}}$ ξ -regular, such that there exists an optimal path π^* with weak δ -clearance for some $\delta > 0$. Let c_N denote the cost of the path returned by DFMT* with N vertices using the cost threshold*

$$r_N = (1 + \eta)^{1/\tilde{D}} C_{\Sigma, \mathcal{M}_{\text{free}}} \left(\frac{\log N}{N} \right)^{1/\tilde{D}},$$

where $\tilde{D} = (n + D)/2$, $C_{\Sigma, \mathcal{M}_{\text{free}}} = \left(C_\mu^{-1} \tilde{D}^{-1} 6^{n+D/2} 2^{n/2} \mu[\mathcal{M}_{\text{free}}] \right)^{1/\tilde{D}}$, and $\eta \geq 0$ is an implementation-specific parameter. Then for fixed $\varepsilon > 0$,

$$\mathbb{P}[c_N > (1 + \varepsilon)c(\pi)] = O\left(N^{-\eta/\tilde{D}} \log^{-1/\tilde{D}} N\right).$$

VII. NUMERICAL EXPERIMENTS

The DFMT* and DPRM* algorithms were implemented in Julia and run using a Unix operating system with a 2.0 GHz processor and 8 GB of RAM. We tested DFMT* and DPRM* on the double integrator system, a standard LQDMP formulation as studied in [9]. We also implemented variants of DFMT* and DPRM* where the local connection cost is computed with respect to a fixed time τ , instead of optimizing $c^*[\tau]$ over all arrival times. This is less computationally intensive than searching for the optimal τ^* , and can be proven asymptotically optimal as well using a similar probabilistic exhaustivity approach. The intuition is that for a fixed cost radius, the algorithm is searching over a “donut” – the band of states in which the fixed τ is a valid approximation for τ^* – instead of a “ball.” Provided the sample count N is sufficiently large, this donut has enough volume to support the desired tracing connections. We denote the optimal-connection algorithms presented in this paper τ^* DFMT* and τ^* DPRM* to differentiate the two approaches. The simulation results are summarized in Figure 1. A maze was used for \mathcal{M}_{obs} , and our algorithm implementations were run 50 times each on sample sizes up to $N = 12000$ for fixed τ and $N = 6000$ for the τ^* variants. We plot results for both versions of DFMT* run with and without a cache of near neighbor sets and local connection costs; as discussed in [10] this information, which does not depend on the problem-specific obstacle configuration, may be precomputed for batch-processing algorithms such

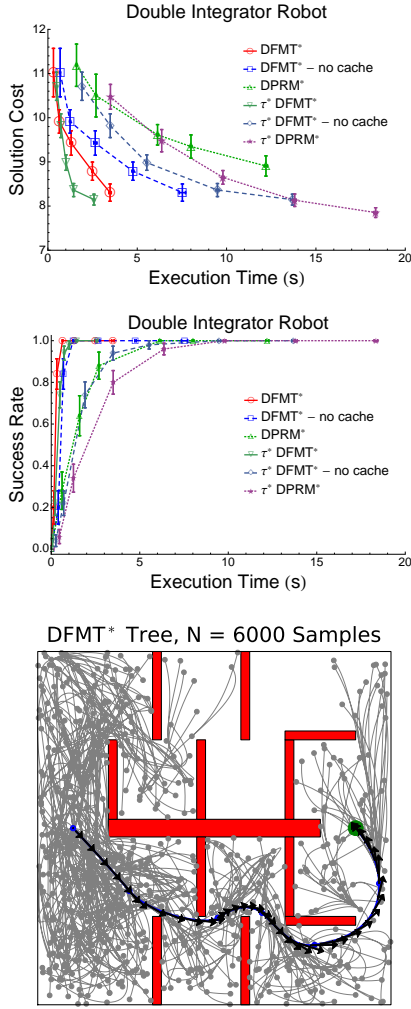


Fig. 1. Top: Simulation results for the double integrator system with a maze obstacle set. The error bars in each axis represent plus and minus one standard error of the mean for a fixed sample size n . Bottom: Example DFMT* tree for $n = 6000$. The position of the feasible trajectory returned is highlighted in blue with velocity denoted by arrows.

as DFMT* and DPRM* – the price to pay is a moderate increase in memory requirements. We see that the extra time for optimizing over local connection duration τ is significant (DFMT* – no cache vs. τ^* DFMT* – no cache), but may be mitigated by precomputation (DFMT* vs. τ^* DFMT*).

VIII. DISCUSSION AND CONCLUSIONS

In this paper we have provided a thorough and rigorous theoretical framework to assess optimality guarantees of sampling-based algorithms for linear affine systems with a mixed time/energy cost function. In particular, we leveraged the study of small-cost perturbations to show that optimum-approximating waypoints may be found among randomly sampled state sets with high probability. We applied this analysis to design and theoretically validate an asymptotically optimal algorithm, DFMT*, for the LQDMP problem.

Although this work is nominally limited to linear affine drift systems, it not only provides a good model for many real systems, but is a crucial first step towards modelling nonlinear systems as well. Indeed, since DFMT* can be applied to a nonlinear system by linearizing the dynamics, an important

next step will be to assess the theoretical guarantees of DFMT* applied to such a linearized approximation. Given suitable smoothness assumptions on the system linearization (sufficiently powerful, but still encompassing a useful class of dynamics), it seems likely that the perturbation analysis and probabilistic exhaustivity may follow identically to that presented in this paper, up to an additional term quantifying the “perturbation on the perturbation.” A similar linearization approach has been experimentally validated by Kinodynamic RRT* in [9]. We also note the similarities in analysis evident between the linear affine drift systems studied in this paper and the (possibly non-linear) control-affine driftless systems studied in [10]. In particular, the parallel notions of controllable/bracket-generating systems and controllability index/Hausdorff dimension give hope that a unifying theory for non-linear systems with drift may be achieved. There are a number of additional directions open for further research. In particular, we plan to deploy DFMT* on robotic platforms, specifically helicopters and floating platforms emulating the dynamics of spacecraft. Also, it is of interest to study a bidirectional version of DFMT*. Finally, it is of interest to devise strategies whereby the radius tuning parameter is self regulating, with the objective of making the algorithm parameter-free.

REFERENCES

- [1] S. M. LaValle, *Planning Algorithms*. Cambridge University Press, 2006.
- [2] —, “Motion planning: Wild frontiers,” *IEEE Robotics Automation Magazine*, vol. 18, no. 2, pp. 108–118, 2011.
- [3] S. M. LaValle and J. J. Kuffner, “Randomized kinodynamic planning,” *International Journal of Robotics Research*, vol. 20, no. 5, pp. 378–400, 2001.
- [4] S. Karaman and E. Frazzoli, “Sampling-based algorithms for optimal motion planning,” *International Journal of Robotics Research*, vol. 30, no. 7, pp. 846–894, 2011.
- [5] —, “Optimal kinodynamic motion planning using incremental sampling-based methods,” in *Proc. IEEE Conf. on Decision and Control*, 2010, pp. 7681–7687.
- [6] A. Perez, R. Platt, G. Konidaris, L. Kaelbling, and T. Lozano-Perez, “LQR-RRT*: Optimal sampling-based motion planning with automatically derived extension heuristics,” in *Proc. IEEE Conf. on Robotics and Automation*, 2012, pp. 2537–2542.
- [7] G. Goretin, A. Perez, R. Platt Jr., and G. Konidaris, “Optimal sampling-based planning for linear-quadratic kinodynamic systems,” in *Proc. IEEE Conf. on Robotics and Automation*, 2013, pp. 2429–2436.
- [8] S. Karaman and E. Frazzoli, “Sampling-based optimal motion planning for non-holonomic dynamical systems,” in *Proc. IEEE Conf. on Robotics and Automation*, 2013, pp. 5041–5047.
- [9] D. J. Webb and J. van den Berg, “Kinodynamic RRT*: Optimal motion planning for systems with linear differential constraints,” in *Proc. IEEE Conf. on Robotics and Automation*, 2013, pp. 5054–5061.
- [10] E. Schmerling, L. Janson, and M. Pavone, “Optimal sampling-based motion planning under differential constraints: the driftless case,” in *Proc. IEEE Conf. on Robotics and Automation*, 2015, pp. 2368–2375.
- [11] Y. Li, Z. Littlefield, and K. Bekris, “Sparse methods for efficient asymptotically optimal kinodynamic planning,” *Workshop on Algorithmic Foundations of Robotics*, 2014.
- [12] C.-T. Chen, *Linear system theory and design*. Oxford University Press, 1995.
- [13] F. L. Lewis and V. L. Syrmos, *Optimal control*. New York: Wiley, 1995.
- [14] T. Kailath, *Linear Systems*. Englewood Cliffs, New Jersey: Prentice-Hall, 1980.
- [15] L. Janson, E. Schmerling, A. Clark, and M. Pavone, “Fast marching tree: A fast marching sampling-based method for optimal motion planning in many dimensions,” *International Journal of Robotics Research*, vol. 34, no. 7, pp. 883–921, 2015.